

50325-0796 (Seq. No. 7555)

Patent

UNITED STATES PATENT APPLICATION

FOR

LINK SIZING BASED ON BOTH USER BEHAVIOR AND TRAFFIC CHARACTERISTICS

INVENTORS:

MICHAEL CHEUNG
MARK AMMAR RAYES

PREPARED BY:

HICKMAN PALERMO TRUONG & BECKER LLP
1600 WILLOW STREET
SAN JOSE, CA 95125
(408) 414-1080

EXPRESS MAIL MAILING INFORMATION

"Express Mail" mailing label number EV322192416US

Date of Deposit August 21, 2003

LINK SIZING BASED ON BOTH USER BEHAVIOR AND TRAFFIC CHARACTERISTICS

FIELD OF THE INVENTION

[0001] The present invention generally relates to computer network analysis. The invention relates more specifically to a method and apparatus for link sizing based on both user behavior and traffic characteristics.

BACKGROUND OF THE INVENTION

[0002] The approaches described in this section could be pursued, but are not necessarily approaches that previously have been conceived or pursued. Therefore, unless otherwise indicated herein, the approaches described in this section are not prior art to the claims in this application and are not admitted to be prior art by inclusion in this section.

[0003] Computers and other devices can be enabled to communicate with each other through a computer network, such as a local area network (LAN), wide area network (WAN), or inter-network. In some networks, end stations, such as computers, are connected to intermediate network elements, such as routers, through communication links. The intermediate network elements also are connected to each other through communication links. Thus, one or more paths of communication links are established between end stations in a network.

[0004] A communication link, such as a cable, has a capacity. The capacity of a communication link describes the rate at which the communication link can transmit data. The rate at which data can be transmitted is often called “bandwidth.” For example, a communication link’s bandwidth may be expressed as a number of bits per second. Some kinds of communication links have greater capacities than other kinds of communication links. The capacity of a communication link may be called the “size” of the communication link.

[0005] A single communication link may carry data transmitted to and from multiple end stations. For example, multiple end stations may be connected to single intermediate network element, which may be connected to a network through a single communication link. In such a configuration, all communication between any of the end stations and the network passes through the single communication link. Consequently, all of the end stations share the bandwidth of the single communication link. While a user of one end station is using a portion of the bandwidth, that portion is unavailable to users of the other end stations.

[0006] Users may enter into agreements with network service providers in order to obtain network connectivity. Such agreements are often called “service level agreements,” or “SLAs.” For example, a user may agree to pay a network service provider a specified amount of money per month in exchange for network access. Often, an agreement guarantees a minimum quality of service (QoS) and/or grade of service (GoS) to a user.

[0007] A QoS may specify several factors, such as the maximum probability that a data packet will be lost, and/or the maximum delay that a data packet will experience in transit. Generally, QoS factors relate to characteristics of data packets. In other words, QoS factors generally relate to traffic characteristics.

[0008] A GoS also may specify several factors, such as the maximum probability that a user’s attempt to establish a communication session with the network will fail. A user’s attempt to establish a communication session with the network may fail if a communication link’s available bandwidth is less than the bandwidth that the communication session requires. In a converged network that provides packet voice service, a communication session is referred to as a “call.”

[0009] For example, when a user dials a telephone number on a telephone, the user initiates a call. When the user hangs up the telephone, the user terminates the call. For another example,

when a user logs on to a network, the user initiates a call. When the user logs off of the network, the user terminates the call. The probability that an attempt to establish a communication session will fail is referred to as the “call blocking probability.” Thus, a GoS may specify a maximum call blocking probability. Generally, GoS factors relate to characteristics of calls.

[0010] As mentioned above, multiple end stations may share the bandwidth of a single communication link. Each such end station may be associated with a user who has been guaranteed a minimum QoS and a minimum GoS. The number of users that share a single communication link may fluctuate. For example, additional users may subscribe to the services of a network server provider that provides network access through the single communication link. Furthermore, the amount of bandwidth used by each user may fluctuate. When the number of users that share a current communication link increases, or when the combined amount of bandwidth used by the users increases, the bandwidth of the current communication link might become less than that necessary to guarantee the users the minimum QoS and the minimum GoS. In order to maintain this guarantee, the current communication link might need to be replaced with a communication link that has greater bandwidth.

[0011] Typically, bandwidth comes at a price. Communication links that provide greater bandwidth also cost more. For example, a T3 link provides greater bandwidth than a T1 link, but also costs more than the T1 link. To reduce costs, it is desirable to determine the minimum increase in bandwidth that is needed to support a specified increase in the number of users that share a communication link. The process of making this determination may be referred to as “link sizing.” Some previous approaches to making this determination exist.

[0012] Some approaches make the determination based solely on call characteristics. Call characteristics are a product of user behavior. For example, some approaches make the determination based solely on the average time between the arrivals of new calls on a

communication link (the “inter-call arrival time”) and the average duration of calls on a communication link. Such approaches typically are used to determine the capacities of links in telephone networks. For each call currently on a communication link in a telephone network, a separate portion of the capacity of the communication link is dedicated exclusively to that call. Regardless of whether participants in a particular call are actively using the portion dedicated to the call, for example, by speaking, the portion is not made available to other calls during the particular call.

[0013] These call-characteristic-only-based approaches typically overestimate needed bandwidth when applied to networks in which a portion of overall bandwidth is not dedicated exclusively to a call. In Internet Protocol (IP) networks, the overall bandwidth of a communication link may be allocated among calls based on the calls’ usage of the bandwidth at a given moment. For example, if a first call does not actively transmit data packets during a period of time, then, during that period of time, a second call may transmit data packets using the bandwidth formerly used by the first call. When the first call again actively transmits data packets, the second call may cease using the bandwidth being used by the first call. In this manner, less available bandwidth is wasted, and a greater number of calls may be supported. However, because call-characteristic-only-based approaches do not account for the sharing of bandwidth in this manner, call-characteristic-only-based approaches may determine an amount greater than the minimum amount of bandwidth that actually is needed to sustain a specified number of users.

[0014] Other approaches determine the needed increase in bandwidth based solely on traffic characteristics. For example, some approaches make the determination based solely on data burst characteristics during calls and how frequently data packets arrive during a burst on a communication link. A burst period is a period of time during which data packets arrive on a

communication link relatively continuously. Such approaches typically are used to determine the capacities of links in IP networks. Examples of approaches that are based solely on traffic characteristics are discussed in A.I. Elwalid and D. Mitra, "Effective Bandwidth of General Markovian Traffic Sources and Admission Control of High Speed Networks," *IEEE/ACM Transactions on Networking*, vol. 1, pp. 329-343, June 1993; D. Anick, D. Mitra, and M. M. Sondhi, "Stochastic Theory of a Data-Handling System with Multiple Sources," *The Bell System Technical Journal*, vol. 51, pp. 1871-1894, October 1982; and M. Schwartz, *Broadband Integrated Networks*, Prentice Hall PTR, 1996.

[0015] These traffic-characteristic-only-based approaches typically overestimate needed bandwidth when applied to networks in which calls are discrete and of limited duration. In a telephone network, a particular user might establish a call that lasts for a limited duration, and then terminate the call. A substantial amount of time might pass before the particular user establishes another call. Because traffic-characteristic-only-based approaches do not account for the limited duration of calls in some networks, traffic-characteristic-only-based approaches may determine an amount greater than the minimum amount of bandwidth that actually is needed to sustain a specified number of users. Some of the shortcomings of traffic-characteristic-only-based approaches are discussed in M. Beshai, R. Kositpaiboon, and J. Yan, "Interaction of Call Blocking and Cell Loss in an ATM Network," *IEEE Journal on Selected Areas of Communications*, vol. 12, pp. 1051-1058, August 1994.

[0016] The shortcomings of the approaches described above are at least partially a consequence of the limited scope of information that each such approach considers in making a determination of needed bandwidth. Based on the foregoing, there is a clear need for a method that considers a broader scope of information in determining the minimum amount of bandwidth needed on a communication link.

BRIEF DESCRIPTION OF THE DRAWINGS

[0017] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

[0018] FIG. 1 is a flow diagram that illustrates a high level overview of one embodiment of a method for link sizing based on both user behavior and traffic characteristics;

[0019] FIG. 2 is a flow diagram that illustrates one embodiment of a method for link sizing based on both user behavior and traffic characteristics;

[0020] FIG. 3 is a flow diagram that illustrates one embodiment of a method for solving a constraint optimization problem when traffic sources and user calling behavior are modeled by a two state on/off process and a Poisson process, respectively; and

[0021] FIG. 4 is a block diagram that illustrates a computer system upon which an embodiment may be implemented.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

[0022] A method and apparatus for link sizing based on both user behavior and traffic characteristics is described. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

[0023] Embodiments are described herein according to the following outline:

- 1.0 General Overview
- 2.0 Functional Overview
- 3.0 Method of Link Sizing Based on Both User Behavior and Traffic Characteristics
 - 3.1 Formulating and Solving the Constraint Optimization Problem
 - 3.2 Modeling User Calling Behavior and Traffic Sources Using a Poisson Process and a Two State On/Off Process
- 4.0 Implementation Mechanisms—Hardware Overview
- 5.0 Extensions and Alternatives

1.0 GENERAL OVERVIEW

[0024] The needs identified in the foregoing Background, and other needs and objects that will become apparent from the following description, are achieved in the present invention, which comprises, in one aspect, a method for link sizing based on both user behavior and traffic characteristics. According to one aspect of the method, an amount of bandwidth needed on a

link (i.e., a link size) is determined based on both user behavior and traffic characteristics. The determined amount is stored in memory.

[0025] Unlike other approaches to link sizing, techniques disclosed herein do not restrict the information considered in determining the link size to user behavior alone or traffic behavior alone. As a result, techniques disclosed herein determine a minimum required link size with greater accuracy than other approaches.

[0026] In other aspects, the invention encompasses a computer apparatus and a computer-readable medium configured to carry out the foregoing steps.

2.0 FUNCTIONAL OVERVIEW

[0027] FIG. 1 is a flow diagram that illustrates a high level overview of one embodiment of a method 100 for link sizing based on both user behavior and traffic characteristics. Such a method may be performed by any of many different mechanisms, such as, for example, computer system 400 described below with reference to FIG. 4, or a network management software application that is hosted by system 400.

[0028] In block 102, a specified number of users to be supported by a communication link is received. For example, the specified number of users may be the maximum number of users that might simultaneously attempt to participate in calls over the communication link.

[0029] In block 104, one or more specified GoS factors are received. For example, the one or more GoS factors may comprise a maximum call blocking probability requirement that is specified in a Service Level Agreement (SLA). A maximum call blocking probability requirement indicates the maximum acceptable probability that an attempt to establish a call will fail because of insufficient bandwidth on the communication link.

[0030] In block 106, one or more specified QoS factors are received. For example, the one or more QoS factors may comprise a specified maximum packet loss probability requirement. For another example, the one or more QoS factors may comprise a specified maximum packet delay requirement. A maximum packet loss probability requirement indicates the maximum acceptable probability that a data packet transmitted on one end of the communication link will not be received on the other end of the communication link. A maximum packet delay requirement indicates the maximum acceptable amount of time that may occur between a data packet being transmitted on one end of the communication link and the data packet being received on the other end of the communication link. The data received in blocks 102, 104, and 106 may be received using any computer-implemented input process, such as user input to a software-generated user interface, reading from a data file, receiving a parameter value in a function call, method invocation, etc.

[0031] In block 108, user behavior relative to the communication link is determined. For example, the behavior of users who used the communication link in the past may be determined based on records of such behavior. User behavior may be determined based on the characteristics of calls made over the communication link. As described above, call characteristics may comprise the average inter-call arrival time and the average call duration. Call detail records (“CDRs”) managed by a call manager may be used as source data.

[0032] In block 110, characteristics of traffic on the communication link are determined. For example, the characteristics of past traffic on the communication link may be determined based on records of such characteristics. As described above, traffic characteristics may comprise the frequency with which data packets arrive during calls on the communication link. Traffic characteristics also may comprise the average duration of burst periods during calls on the communication link.

[0033] In block 112, based on both the user behavior and the traffic characteristics, the minimum amount of bandwidth required for the communication link to support the specified number of users while satisfying the specified QoS and GoS factors is determined. The minimum amount may be approximated.

[0034] In block 114, the determined minimum amount of bandwidth is stored in memory. For example, the minimum amount may be stored in main memory 406 of computer system 400 illustrated in FIG. 4. The minimum amount may be presented to a user through a mechanism such as display 412, or communicated, using any programmatic data communication mechanism or messaging mechanism, to another system, application, or process. Based on the minimum amount, the user may replace an existing communication link with a communication link having the minimum amount of bandwidth determined.

[0035] By accounting for user behavior and traffic characteristics in combination rather than only one or the other in isolation, the required size of a communication link may be estimated more accurately. As a result, the cost associated with upgrading an existing communication link is minimized.

3.0 METHOD OF LINK SIZING BASED ON BOTH USER BEHAVIOR AND TRAFFIC CHARACTERISTICS

3.1 FORMULATING AND SOLVING THE CONSTRAINT OPTIMIZATION PROBLEM

[0036] The link-sizing problem addressed in block 112 and elsewhere herein may be expressed as a constraint optimization problem. In this approach, given a current number of users K that use a current communication link with a capacity S , a minimum capacity S' of a communication link is determined subject to the constraints that:

$$\sum_{i=0}^{K'} P_i \cdot \varepsilon_i(S') \leq \theta$$

and:

$$S \leq S'$$

where K' is a specified number of users that the communication link will need to support, K' is typically greater than K , P_i is the probability that i users are actively using the communication link while satisfying a specified maximum call blocking probability, $\varepsilon_i(S')$ is the packet loss probability when i users are using a communication link having capacity S' , and θ is a specified maximum packet loss probability. The specified GoS factor, which in this case is the specified maximum call blocking probability, is captured in the determination of P_i . The specified QoS factor, which in this case is the specified maximum packet loss probability, is θ .

[0037] The constraint optimization problem considers both user behavior and traffic characteristics. User behavior is expressed through P_i , while traffic characteristics are expressed through $\varepsilon_i(S')$. The combination of user behavior and traffic characteristics is expressed through the product of these terms, $P_i \cdot \varepsilon_i(S')$. This is allowed because the user behavior and the traffic characteristics are independent. That is, the user calling characteristics are independent of the traffic being generated during a call. Therefore, the determination of S' , which is accomplished by solving the constraint optimization problem, takes into account a combination of user behavior and traffic characteristics.

[0038] FIG. 2 is a flow diagram that illustrates one embodiment of a method 200 for solving a constraint optimization problem that accounts for a combination of user behavior and traffic characteristics. Such a method may be performed by any of many different mechanisms, such

as, for example, computer system 400 described below with reference to FIG. 4, or an application hosted by system 400.

[0039] In block 202, for each i from 0 to K' , inclusive, a separate corresponding P_i is determined. P_i is the probability that i users are concurrently and actively using the communication link if the specified maximum call blocking probability is satisfied. Each P_i may be referred to as a “state probability.” Each state probability may be determined, for example, according to the well-known Erlang Loss Formula. Other techniques for determining each state probability also may be used.

[0040] In block 204, the initial required capacity S' is guessed. For example, S' initially may be guessed to be S .

[0041] In block 206, for each i from 0 to K' , inclusive, a separate corresponding $\epsilon_i(S')$ is determined. $\epsilon_i(S')$ is probability that a given data packet will be lost when i users are using a communication link having capacity S' . Each $\epsilon_i(S')$ may be referred to as a “marginal packet loss probability.” Each marginal packet loss probability may be determined, for example, according to the well-known Fluid-Flow Analysis. Techniques for determining marginal packet loss probabilities are described in D. Anick, D. Mitra, and M. M. Sondhi, “Stochastic Theory of a Data-Handling System with Multiple Sources,” *The Bell System Technical Journal*, vol. 51, pp. 1871-1894, October 1982; and M. Schwartz, *Broadband Integrated Networks*, Prentice Hall PTR, 1996. Other techniques for determining each marginal packet loss probability also may be used.

[0042] In block 208, an overall packet loss rate is determined by adding the products of P_i and $\epsilon_i(S')$ for each i from 0 to K' , inclusive. Multiplying the state probability with the corresponding marginal packet loss probability may be referred to as “weighting” the marginal packet loss

probability with the corresponding state probability. The summation of the weighted probabilities is the overall packet loss rate. In other words, the overall packet loss rate on a communication link having a capacity S' may be expressed as:

$$\sum_{i=0}^{K'} P_i \bullet \epsilon_i(S').$$

[0043] In block 210, it is determined whether the overall packet loss rate satisfies the specified maximum packet loss probability θ . The overall packet loss rate satisfies the specified maximum packet loss probability if the overall packet loss rate is less than or equal to the specified maximum packet loss probability. If the overall packet loss rate satisfies the specified maximum packet loss probability, then control passes to block 214. Otherwise, control passes to block 212.

[0044] In block 212, the capacity S' is incremented by a specified amount ΔS . For example, ΔS may be the difference between the bandwidths of two commercially available communication links. Thus, if S' is the bandwidth of a T1 link, then ΔS may be selected to be the difference between the bandwidth of a T1 link and the bandwidth of a T3 link. Each time that S' is incremented in this manner, ΔS may be a different value. After S' has been incremented, control passes back to block 206.

[0045] Alternatively, in block 214, the capacity S' is stored in memory as the minimum capacity needed on a communication link that is to support K' users. S' may be presented to a user.

[0046] By incrementally determining whether various values of S' will cause the overall packet loss rate to satisfy the specified maximum packet loss probability concurrently with the satisfaction of the specified maximum call blocking probability, the minimum required capacity of a communication link needed to support a specified number of users may be determined.

3.2 MODELLING USER CALLING BEHAVIOR AND TRAFFIC SOURCES USING A POISSON PROCESS AND A TWO STATE ON/OFF PROCESS

[0047] Many different techniques may be used to model user calling behavior and traffic sources. The selection of the techniques used to determine state probabilities and marginal packet loss probabilities may be based on the selected user calling behavior and traffic source models.

[0048] A traffic source may be modeled by a two state on/off process, with on and off periods being exponentially distributed with means α^{-1} and η^{-1} , respectively, and the packet arrival rate during the on period, γ . On periods are burst periods. Off periods are idle periods during which packets are not transmitted during a call. Voice traffic may be modeled in this way. The lengths of on and off periods during a call are independent of the number of calls made during a period of time. Thus, traffic characteristics are independent of user calling behavior.

[0049] Considering identical traffic sources and a communication link using a single-server queue with infinite waiting room, Fluid Flow Analysis may be used to estimate the packet loss rate by the tail distribution of the buffer occupancy l when the buffer occupancy exceeds a prescribed buffer size B , i.e., $P(l < B)$. The packet loss rate may be determined by solving a set of differential equations that describe the equilibrium buffer occupancy. Such a set is described in D. Anick, D. Mitra, and M. M. Sondhi, "Stochastic Theory of a Data-Handling System with Multiple Sources," *The Bell System Technical Journal*, vol. 51, pp. 1871-1894, October 1982; and M. Schwartz, *Broadband Integrated Networks*, Prentice Hall PTR, 1996.

[0050] User calling behavior may be modeled by a Poisson process, such that average inter-call arrival time and average call duration are exponentially distributed with means λ^{-1} and μ^{-1} , respectively. The communication link may be modeled by an M/M/K'/K' queuing system,

which is also known as the Erlang Model. User behavior may be analyzed using the Erlang Loss Formula.

[0051] Thus, according to one embodiment, it is desirable to determine the minimum capacity S' subject to the constraints that:

$$\sum_{i=0}^{K'} P_i \bullet \varepsilon_i(\alpha, \eta, \gamma, B, S') \leq \theta$$

and:

$$S \leq S'$$

where K' is the specified number of users that the communication link will need to support; P_i is the probability that i users are actively using the communication link while satisfying the specified maximum call blocking probability (using the Erlang Model); $\varepsilon_i(\alpha, \eta, \gamma, B, S')$ is the packet loss probability (using the Fluid Flow Model) when i users are using a communication link having capacity S' and an associated buffer size B , when the mean on period is α^{-1} , the mean off period is η^{-1} , and the packet rate during the on period is γ ; and θ is the specified maximum packet loss probability.

[0052] FIG. 3 is a flow diagram that illustrates one embodiment of a method 300 for solving a constraint optimization problem when data traffic and user calling behavior are modeled by a two state on/off process and a Poisson process, respectively. Such a method may be performed by any of many different mechanisms, such as, for example, computer system 400 described below with reference to FIG. 4.

[0053] In block 302, based on a specified average call duration μ^{-1} , a specified number of users K' , and a specified maximum call blocking probability, an average inter-call arrival time λ is determined using M/M/K'/K' Analysis (Erlang Analysis).

[0054] In block 304, for each i from 0 to K' , inclusive, a separate corresponding P_i is determined based on the average inter-call arrival time λ . P_i is the probability that i users are concurrently and actively using the communication link if the specified maximum call blocking probability is satisfied.

[0055] In block 306, the initial required capacity S' is guessed. For example, S' initially may be guessed to be S .

[0056] In block 308, for each i from 0 to K' , inclusive, a separate corresponding $\varepsilon_i(\alpha, \eta, \gamma, B, S')$ is determined using Fluid-Flow Analysis. $\varepsilon_i(\alpha, \eta, \gamma, B, S')$ is probability that a given data packet will be lost when i users are using a communication link having capacity S' and an associated buffer size B , when the mean on period is α^{-1} , the mean off period is η^{-1} , and the packet rate during the on period is γ .

[0057] In block 310, an overall packet loss rate is determined by adding the products of P_i and $\varepsilon_i(\alpha, \eta, \gamma, B, S')$ for each i from 0 to K' , inclusive. The overall packet loss rate on a communication link having a capacity S' and a buffer size B , when the mean on period is α^{-1} , the mean off period is η^{-1} , and the packet rate during the on period is γ , may be expressed as:

$$\sum_{i=0}^{K'} P_i \bullet \varepsilon_i(\alpha, \eta, \gamma, B, S').$$

[0058] In block 312, it is determined whether the overall packet loss rate satisfies the specified maximum packet loss probability θ . The overall packet loss rate satisfies the specified maximum packet loss probability if the overall packet loss rate is less than or equal to the specified maximum packet loss probability. If the overall packet loss rate satisfies the specified maximum packet loss probability, then control passes to block 316. Otherwise, control passes to block 314.

[0059] In block 314, the capacity S' is incremented by a specified amount ΔS . After S' has been incremented, control passes back to block 308.

[0060] Alternatively, in block 316, the capacity S' is stored in memory as the minimum capacity needed on a communication link that is to support K' users. S' may be presented to a user.

4.0 IMPLEMENTATION MECHANISMS -- HARDWARE OVERVIEW

[0061] FIG. 4 is a block diagram that illustrates a computer system 400 upon which an embodiment of the invention may be implemented. Computer system 400 includes a bus 402 or other communication mechanism for communicating information, and a processor 404 coupled with bus 402 for processing information. Computer system 400 also includes a main memory 406, such as a random access memory ("RAM") or other dynamic storage device, coupled to bus 402 for storing information and instructions to be executed by processor 404. Main memory 406 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 404. Computer system 400 further includes a read only memory ("ROM") 408 or other static storage device coupled to bus 402 for storing static information and instructions for processor 404. A storage device 410, such as a magnetic disk or optical disk, is provided and coupled to bus 402 for storing information and instructions.

[0062] Computer system 400 may be coupled via bus 402 to a display 412, such as a cathode ray tube ("CRT"), for displaying information to a computer user. An input device 414, including alphanumeric and other keys, is coupled to bus 402 for communicating information and command selections to processor 404. Another type of user input device is cursor control 416, such as a mouse, trackball, stylus, or cursor direction keys for communicating direction

information and command selections to processor 404 and for controlling cursor movement on display 412. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

[0063] The invention is related to the use of computer system 400 for link sizing based on both user behavior and traffic characteristics. According to one embodiment of the invention, link sizing based on both user behavior and traffic characteristics is provided by computer system 400 in response to processor 404 executing one or more sequences of one or more instructions contained in main memory 406. Such instructions may be read into main memory 406 from another computer-readable medium, such as storage device 410. Execution of the sequences of instructions contained in main memory 406 causes processor 404 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

[0064] The term “computer-readable medium” as used herein refers to any medium that participates in providing instructions to processor 404 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 410. Volatile media includes dynamic memory, such as main memory 406.

Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 402. Transmission media can also take the form of acoustic or light waves, such as those generated during radio wave and infrared data communications.

[0065] Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other

optical medium, punchcards, papertape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

[0066] Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to processor 404 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 400 can receive the data on the telephone line and use an infrared transmitter to convert the data to an infrared signal. An infrared detector can receive the data carried in the infrared signal and appropriate circuitry can place the data on bus 402. Bus 402 carries the data to main memory 406, from which processor 404 retrieves and executes the instructions. The instructions received by main memory 406 may optionally be stored on storage device 410 either before or after execution by processor 404.

[0067] Computer system 400 also includes a communication interface 418 coupled to bus 402. Communication interface 418 provides a two-way data communication coupling to a network link 420 that is connected to a local network 422. For example, communication interface 418 may be an integrated services digital network (“ISDN”) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 418 may be a local area network (“LAN”) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 418 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

[0068] Network link 420 typically provides data communication through one or more networks to other data devices. For example, network link 420 may provide a connection through local network 422 to a host computer 424 or to data equipment operated by an Internet Service Provider (“ISP”) 426. ISP 426 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the “Internet” 428. Local network 422 and Internet 428 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 420 and through communication interface 418, which carry the digital data to and from computer system 400, are exemplary forms of carrier waves transporting the information.

[0069] Computer system 400 can send messages and receive data, including program code, through the network(s), network link 420 and communication interface 418. In the Internet example, a server 430 might transmit a requested code for an application program through Internet 428, ISP 426, local network 422 and communication interface 418. In accordance with the invention, one such downloaded application provides for link sizing based on both user behavior and traffic characteristics as described herein.

[0070] The received code may be executed by processor 404 as it is received, and/or stored in storage device 410, or other non-volatile storage for later execution. In this manner, computer system 400 may obtain application code in the form of a carrier wave.

5.0 EXTENSIONS AND ALTERNATIVES

[0071] In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the

invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

[0072] For example, while a technique described above can be used for “upsizing” a link, a similar technique may be used for “downsizing” a link to accommodate decreased demand. When downsizing a link, the constraint $S \leq S'$ in the optimization problem described above changes to $S \geq S'$, and K typically is greater than K' rather than vice-versa. Instead of incrementing S' by a specified amount, as in blocks 214 and 314 above, S' is decremented by the specified amount. Instead of storing S' as the minimum capacity required when the overall packet loss rate satisfies the specified maximum packet loss probability, as in blocks 210, 214, 312, and 316 above, S' is stored as the minimum capacity required when further decrementing S' by the specified amount would cause the overall packet loss rate to not satisfy the specified maximum packet loss probability.
